
Theta-RBM: Unfactored Gated Restricted Boltzmann Machine for Rotation-Invariant Representations

Mario Valerio Giuffrida
 PRIAn Research Unit
 IMT School Advanced Studies
 Lucca (LU), Italy
 valerio.giuffrida@imtlucca.it

Sotirios A. Tsafaris
 School of Engineering
 University of Edinburgh
 Edinburgh, UK
 s.tsafaris@ed.ac.uk

Abstract

Learning invariant representations is a critical task in computer vision. In this paper, we propose the Theta-Restricted Boltzmann Machine (θ -RBM in short), which builds upon the original RBM formulation and injects the notion of rotation-invariance during the learning procedure. In contrast to previous approaches, we do not transform the training set with all possible rotations. Instead, we rotate the gradient filters when they are computed during the *Contrastive Divergence* algorithm. We formulate our model as an unfactored gated Boltzmann machine, where another input layer is used to modulate the input visible layer to drive the optimisation procedure. Among our contributions is a mathematical proof that demonstrates that θ -RBM is able to learn rotation-invariant features according to a recently proposed invariance measure. Our method reaches an invariance score of $\sim 90\%$ on *mnist-rot* dataset, which is the highest result compared with the baseline methods and the current state of the art in transformation-invariant feature learning in RBM. Using an SVM classifier, we also showed that our network learns discriminative features as well, obtaining $\sim 10\%$ of testing error.

1 Introduction

Most of the applications in computer vision require suitable image representations, which are invariant to certain geometrical transformations. Recently, learned image representations have demonstrated impressive performance and have become a powerful tool in computer vision. In particular, neural networks have been shown to learn more discriminative image representations, by learning task-specific filters to apply to input images [3]. A well-known neural network that learns features in an unsupervised fashion is the Restricted Boltzmann Machine (RBM) [19], characterised by a bipartite graph whose sides are referred as *visible* and *hidden* units respectively. In its standard formulation RBM does not accommodate nuisance factors in the scene, hence it cannot learn invariant features.

The easiest way to achieve a notion of invariance is to present to the learning algorithm as much as variability as possible. When this is not possible, simulated variability is introduced by applying transformations artificially to data, in what is known as *dataset augmentation*. The act of altering input data has been exploited thoroughly in this regard (e.g. [17, 20]). However, this has disadvantages such as: (i) introducing alias due to pixel interpolation in the transformed images; and (ii) the transformed data may not span over the entire space of transformations. Instead, recently in [1], gradient information is exploited to transform data into a common reference frame and thus eschew augmentation.

In this paper we present the θ -RBM that learns image representations that are invariant to rotations. Taking inspiration by previous works [1, 21], we propose an unfactored gated Restricted Boltzmann Machine that optimises a third-order tensor, which learns rotation-invariant features from input




Figure 1: Representation of our θ -RBM as an unfactored gated Boltzmann machine. x and r are visible layers respectively representing input image patterns and rotations and h is the hidden layer that learns the relationship among inputs via the third-order W shown in the middle.

images. In our network, shown in Figure 1, we use the visible layer x to provide images, while the input layer r is *only* used to indicate the dominant rotation of input patches. An input image is paired with the corresponding rotation, which is found by computing the dominant gradient of this image.

Contributions: Our contributions are: (i) a compact gated formulation of RBMs for rotation invariance; (ii) which mathematically prove that it learns rotation-invariant features, on the basis of an known invariance measure [16]; (iii) we do not require dataset augmentation or pre-training to learn rotated patterns as required by others [21], thus training is faster; (iv) the *unfactored* third-order tensor uses a limited amount of slices leading to a more compact data representation compared with other methods that use several factors [13]; and finally (v) we also show empirically that our method learns rotation-invariant features even when not presented with rotated data (ie. using an unrotated training set).

The remainder of this paper is organised as follows. Section 2 reviews related work. Section 3 presents the proposed unfactored gated formulation of θ -RBM. In Section 4 we provide the mathematical proof of invariant feature learning. In Section 5 we report experimental results of our method on *mnist-rot*. Section 6 offers conclusions.

2 Related Works

Invariance in RBM: Mainly due to their wide utilisation in deep learning, several adaptations and extensions of the original RBM model have been proposed to accommodate transformed variants of the same input. A Deep Belief Network (DBN) [8], obtained by stacking RBMs, can learn higher-level representations with invariance characteristics [6]. In [18] the authors extend the DBN and train it several times, such that at each iteration the learned weights are transformed w.r.t to a set \mathcal{S} of transformations. Such approach increases the computational burden of training a single time a DBN linearly with the cardinality of \mathcal{S} , which may require a considerable amount of time. The Transformation-Invariant RBM (TI-RBM) [20] can learn features that are invariant to local transformations, by altering the training w.r.t. a set of suitable transformations (e.g., rotations, translations, scaling). The main drawbacks of this methodology is very similar to the straightforward approach of dataset augmentation, since transformed images may not fully cover the entire manifold where the data may lie. In [1], patches are transformed beforehand using SIFT features, applying rotation and scaling to normalise the inputs. In [9], the authors present the Spatial Transformer, a differentiable layer incorporated inside a convolutional architecture to learn affine transformations

from the feature maps. Estimating the transformations that apply to an input image showed to be an interesting approach towards discriminative invariant features, and it is the direction we pursued and incorporated on our θ -RBM.

Gated networks: Recently, a radically different view of network architectures has been proposed, known as gated networks [12] which have demonstrated impressive results on learning image pairs. The basic idea of gated networks is to optimise the parameters involved in the model, such that the input layer x is *modulated* by another visible layer y to find a common representation h . In [22], the authors show that a Gated Boltzmann Machine (GBM) can learn rotated versions of faces. Even though the discriminative power of GBMs has been demonstrated in several applications [10, 12, 14, 21, 22], they do need a pair of related images to be trained. Nonetheless, in [21], the authors showed that GBM can be used for a different purpose by just slightly varying the original formulation. In fact, instead using the y layer to provide another image, they use it as feature selection activation layer, such that the model learns which features are more discriminative for the task at hand. These GBMs are able to efficiently use the three-way interaction of two input layers and a hidden layer which in turn provides a common representation for the inputs. Furthermore, in [15] the authors showed that GBM can be extended such that they can deal with four-way interaction, adding a further layer to provide class labels. The two visible layers appearing in a three-way GBM are used during the optimisation to find correlation between them. As a result, the model is able to learn features that are invariant to affine transformations, when pairs of related data are provided (e.g., pairs of faces [22]). Following a similar approach of [21], we use the third layer as input, in order to encode the magnitude of rotation of each input image. Therefore, the learning procedure is strongly conditioned by this interaction, and we rotate the learned filters such that the network learns rotation-invariant features.

Invariance measures: in [6] a measure invariance is introduced, which computes the mean activation of features extracted from transformed images over the mean activation of the features obtained from the original dataset. Even though this metric provides a score of invariance, it is unbounded and therefore it is not easy to understand when a network is reaching the maximum invariance. Another metric that has been utilised in literature is the mean squared error over the L_2 distance of normalised features [23]. While being bounded its main drawback is that vectors need to be normalised prior to computing the invariance score. In fact, magnitude difference, resulting from transformed images, could be due to lack of invariance, which may bias a classification algorithm. In fact in [5], the authors propose to compute the average over the classification scores of the transformed images. The latter approach requires two steps, namely feature learning and classifier training, which increase the whole complexity of the process. Another metric of invariance was proposed in [16], which behaves similar to the autocorrelation. It is bounded in $[0, 1]$ with ‘1’ reflecting full invariance. Amongst the proposed metrics in literature, we chose to develop our theory upon the latter one, because of its defined range and also its computational efficiency when used in empirical demonstrations.

3 θ -Restricted Boltzmann Machine

We present an unfactored gated Boltzmann machine (GBM) for learning rotation-invariant features. The original formulation of unfactored GBM correlates the interaction between the two visible layers and the hidden layer with a third-order tensor $\mathbf{W} \in \mathbb{R}^{H \times V \times S}$ [14], where H is the number of hidden units of the layer h , V is the dimension of the visible layer x , and S is the dimension of the other visible (rotation) layer r , shown pictorially in Figure 1. We build upon this model to propose our θ -RBM. For sake of presentation clarity, the following definitions and analyses are based upon the standard Bernoulli RBM, but the model can be easily extended to treat real-values as well, as described in [2]. In fact, all our experiments are based on a real-valued Gaussian extension.

Firstly, let us define a support set of rotations $\mathcal{S} = \{\varphi_1, \varphi_2, \dots, \varphi_S\}$, containing S rotations. Then, the energy function that characterises our model is

$$E(v, h, r) = \sum_{s=1}^S \sum_{j=1}^H \sum_{k=1}^V r_s (-v_k h_j W_{jks} - b_j h_j - c_k v_k), \quad (1)$$

where \mathbf{v} , \mathbf{h} , and \mathbf{r} are random vectors that can take only binary values. The third-order tensor \mathbf{W} defines the relationship with the layers in the model, c and b define the bias terms for the visible x ,

and hidden layer h respectively. A further constraint we add is the sparseness on \mathbf{r}

$$\sum_{s=1}^S r_s = 1. \quad (2)$$

The vector \mathbf{r} is used as one-hot indicator function, such that if a patch is rotated by φ_s degree, then $r_s = 1$ and the remaining $r_k = 0, \forall k \neq s$. Because \mathcal{S} defines the set of rotations that our model can learn, an input pattern x can be subjected to one rotation. It is easy to check that the conditional probabilities deriving from Equation (1) are

$$p(h_j = 1 | \mathbf{x}, \mathbf{r}) = \sigma \left(\sum_{s=1}^S r_s (b_j + \mathbf{W}_{j, \bullet, s} \mathbf{x}) \right), \quad (3)$$

$$p(x_k = 1 | \mathbf{h}, \mathbf{r}) = \sigma \left(\sum_{s=1}^S r_s (c_k + \mathbf{h}' \mathbf{W}_{\bullet, k, s}) \right), \quad (4)$$

where $\mathbf{W}_{j, \bullet, s}$ denotes the j -th row in the slice s of the third-order tensor \mathbf{W} , $\mathbf{W}_{\bullet, k, s}$ denotes the k -th column in the s -th slice in \mathbf{W} , \mathbf{h}' is the transpose of the column vector \mathbf{h} , and $\sigma(y)$ is the sigmoid function. In our model, the layer r is treated as input and rotation is found as the dominant orientation experimental computed using the angle with highest frequency in the histogram of oriented gradients.

For a single pattern taken from the dataset (\tilde{x}, \tilde{s}) , image patch and orientation respectively, the Contrastive Divergence algorithm [7] is performed as usual between the visible layer x and the hidden layer h . This procedure computes the gradient $\nabla \mathbf{W}_{\tilde{s}}$, that is the \tilde{s} -th frontal slice of the tensor \mathbf{W} (we denoted $\mathbf{W}_{\bullet, \bullet, s}$ as \mathbf{W}_s for brevity). The remaining slices in \mathbf{W} are computed as follows

$$\nabla \mathbf{W}_k = R_\theta(\nabla \mathbf{W}_{\tilde{s}}), \quad \forall k = 1, 2, \dots, S, k \neq \tilde{s}. \quad (5)$$

$$\text{where} \quad R_\theta(\mathbf{A}) = \mathbf{A} \mathbf{T}'_\theta, \quad (6)$$

such that $R_\theta(\mathbf{A})$ is a transformation function that rotates by $\theta = \varphi_t - \varphi_{\tilde{s}}$ degrees each row in \mathbf{A} . The matrix $\mathbf{T} \in \mathbb{R}^{n \times n}$ defines the rotation and is represented as a particular permutation matrix such that $\mathbf{T} \mathbf{T}' = \mathbf{T}' \mathbf{T} = \mathbf{I}$, $\det(\mathbf{T}) = \pm 1$. The input matrix \mathbf{A} is transposed in $R_\theta(\bullet)$ because filters in a slice in \mathbf{W} are disposed row-wise. Specifically, a filter is a single row in any slice of the tensor \mathbf{W} and it has the same dimension of input images. Once the third-order tensor $\nabla \mathbf{W}$ is computed, the weight matrix is updated as follows

$$\mathbf{W}^{(t)} = \mathbf{W}^{(t-1)} + \eta \nabla \mathbf{W}^{(t)} + \alpha^{(t)} \nabla \mathbf{W}^{(t-1)}, \quad (7)$$

where η is the learning rate, $\alpha^{(t)}$ is the momentum at time t . In (7) we use the superscript index t to indicate the iteration time.

4 Proving Rotation Invariance

We want to demonstrate mathematically that the method above learns rotation invariant features. We used the invariance measure proposed in [16]. Henceforth, we will refer to this measure as γ -score and we will provide a general definition, then we will develop our theory upon it. We adopted this measure because it ranges between '0' and '1', which indicate full variance and full invariance to transformations respectively (see also Section 2).

4.1 The γ -score

Let \mathcal{S} be a set of any transformations¹ that can be applied to the dataset \mathcal{X} . Moreover, let $h_j(\mathbf{x})$ be the state of the j -th hidden neuron when the pattern x is provided (e.g., similar to Equation (3)). The mean activation of a given element in the training set \mathcal{X} across all the transformations in \mathcal{S} is computed as follows:

$$\mu_j(\mathbf{x}) = \frac{1}{S} \sum_{s \in \mathcal{S}} h_j(\mathbf{T}_s(\mathbf{x})), \quad (8)$$

¹We slightly abuse the definition of the set here to be any transform just for the purpose of presenting the score but our analysis and results presume that \mathcal{S} contains only rotations.

where S is the cardinality of the support set \mathcal{S} . Hence, the γ -score for the j -th hidden unit is given by

$$\gamma_j = \frac{\text{var} \{ \mu_j(\mathbf{x}) \}_{\mathbf{x} \in \mathcal{X}}}{\text{var} \{ h_j(\mathbf{x}) \}_{\mathbf{x} \in \mathcal{X}}}, \quad (9)$$

that is the variance of the mean activation of all transformed samples in \mathcal{X} over the variance of activations of the original data \mathcal{X} . The invariance measure in Equation (9) achieves its maximum value ‘1’ when features are invariant to the set of transformations \mathcal{S} .

4.2 Theorem statement

In this section, we will prove that our θ -RBM learns rotation-invariant features, when the learning procedure described in Section 3 is used. Since the maximum value of the γ -score is achieved only when numerator and denominator are equal, we will show that the argument of the variance functions μ_j and h_j are equal. For our purposes, the support set of transformations \mathcal{S} is formed by rotations. Before to proceed with the proof, we need the support of the following lemma.

Lemma 1. *Given a third-order tensor $\mathbf{W} \in \mathbb{R}^{H \times V \times S}$ optimised as described in Section 3, then $\mathbf{W}_{s'}^{(t)} = R_\theta(\mathbf{W}_s^{(t)})$, with $\theta = \varphi_{s'} - \varphi_s$, $\varphi_{s'}, \varphi_s \in \mathcal{S}$.*

Proof. We will proceed by induction over the iteration t . For the base case $t = 0$, we impose that:

- $\mathbf{M} \in \mathbb{R}^{H \times V}$ matrix initialised somehow (e.g., normal distribution)²,
- $\mathbf{W}_s^{(0)} = R_{\varphi_s}(\mathbf{M}), \forall s \in \mathcal{S}$.

namely all the slices in $\mathbf{W}^{(0)}$ are initialised as rotated versions of \mathbf{M} , which initially can be any matrix. Now, let us suppose that the lemma is true until $t - 1$, and demonstrate it for t .

$$\begin{aligned} R_\theta(\mathbf{W}_s^{(t)}) &= R_\theta \left(\underbrace{\mathbf{W}_s^{(t-1)} + \eta \nabla \mathbf{W}_s^{(t)} + \alpha^{(t)} \nabla \mathbf{W}_s^{(t-1)}}_{\text{From Equation (7)}} \right) \\ &= R_\theta \left(\mathbf{W}_s^{(t-1)} \right) + \eta R_\theta \left(\nabla \mathbf{W}_s^{(t)} \right) + \alpha^{(t)} R_\theta \left(\nabla \mathbf{W}_s^{(t-1)} \right) \\ &\quad \underbrace{\hspace{10em}}_{\text{Using the linearity of the transformation}} \\ &= \underbrace{\mathbf{W}_{s'}^{(t-1)}}_{\text{By induction}} + \eta \underbrace{\nabla \mathbf{W}_{s'}^{(t)} + \alpha^{(t)} \nabla \mathbf{W}_{s'}^{(t-1)}}_{\text{From Equation (5)}} = \mathbf{W}_{s'}^{(t)}. \end{aligned}$$

■

Theorem 1. *Under the hypotheses of 1 and given a support set \mathcal{S} of S rotations, $\gamma = 1$ for θ -RBM, i.e., features are invariant under the rotations in \mathcal{S} .*

Proof. We have to prove that $\gamma_j = 1, j = 1, 2, \dots, H$. From Equation (9), we will show that the numerator and denominator coincide. Some preliminaries that help exposition. In Equation (6) the rotation transformation was defined as a row-wise operation. Since we used as convention that vectors are column-wise, we will use the following definition

$$R'_\theta(\mathbf{x}) \equiv [R_\theta(\mathbf{x}')]^\top = [\mathbf{x}'^\top \mathbf{T}'_\theta]^\top, \quad (10)$$

²We observed that the base case of the induction can be relaxed. Experimental evidence showed that by initialising \mathbf{W} with random numbers drawn from a normal distribution it is still possible to have rotation-invariant features.




Figure 2: Filters learned from θ -RBM on *mnist-rot* dataset. We show a subset of the learned filters for brevity. Observing the columns, filters are rotated versions of each other.

such that, given a column vector as argument, we obtain a column vector as result. Now, starting from the definition of μ_j showed in Equation (8), we get

$$\begin{aligned}
\mu_j(\mathbf{x}) &= \frac{1}{S} \sum_{\theta \in \mathcal{S}} h_j(R'_\theta(\mathbf{x})) = \frac{1}{S} \sum_{\theta \in \mathcal{S}} \underbrace{\sigma \left(\sum_{t=1}^S r_t (b_j + \mathbf{W}_{j,\bullet,t} R'_\theta(\mathbf{x})) \right)}_{\text{From Equation (3)}} \\
&= \frac{1}{S} \sum_{\theta \in \mathcal{S}} \underbrace{\sigma (b_j + \mathbf{W}_{j,\bullet,s'} R'_\theta(\mathbf{x}))}_{\text{From (2). Let } r_{s'} = 1} = \frac{1}{S} \sum_{\theta \in \mathcal{S}} \underbrace{\sigma (b_j + R_\theta(\mathbf{W}_{j,\bullet,s}) R'_\theta(\mathbf{x}))}_{\text{From Lemma 1, } \exists s : R_\theta(\mathbf{W}_s) = \mathbf{W}_{s'}} \\
&= \frac{1}{S} \sum_{\theta \in \mathcal{S}} \underbrace{\sigma (b_j + \mathbf{W}_{j,\bullet,s} \mathbf{T}'_\theta [\mathbf{x}' \mathbf{T}'_\theta]')}_{\text{From Equations (6) and (10)}} = \frac{1}{S} \sum_{\theta \in \mathcal{S}} \underbrace{\sigma (b_j + \mathbf{W}_{j,\bullet,s} \mathbf{T}'_\theta \mathbf{T}_\theta \mathbf{x})}_{(AB)' = B' A'} \\
&= \frac{1}{S} \sum_{\theta \in \mathcal{S}} \underbrace{\sigma (b_j + \mathbf{W}_{j,\bullet,s} \mathbf{x})}_{\mathbf{T}_\theta \mathbf{T}'_\theta = \mathbf{T}'_\theta \mathbf{T}_\theta = \mathbf{I}} = \frac{1}{S} \sum_{\theta \in \mathcal{S}} \sigma \left(\sum_{t=1}^S r_t (b_j + \mathbf{W}_{j,\bullet,t} \mathbf{x}) \right) \\
&= \frac{1}{S} \sum_{\theta \in \mathcal{S}} h_j(\mathbf{x}) = h_j(\mathbf{x}),
\end{aligned}$$

where $\theta = \varphi_s - \varphi_{s'}$, $\varphi_s, \varphi_{s'} \in \mathcal{S}$. Since $\mu_j(\mathbf{x}) = h_j(\mathbf{x})$, then also their variance over all the samples in the training set is equal. Therefore $\gamma_j = 1$. ■

With this theorem, we mathematically proved that the features learnt by our θ -RBM are invariant under a set of rotations \mathcal{S} .

5 Experimental Results

In this section we show experimentally that θ -RBM is able to learn rotation-invariant features. Briefly summarized, training with *mnist-rot* dataset [11] we will show the claim of 1 and we will estimate invariance according to Equation (9). As baselines, we compared with the standard (Gaussian-Bernoulli) RBM, the *Oriented RBM* (O-RBM), which is trained by pre-aligning the input images by their respectively dominant orientation. We compare also with TI-RBM –a recent state-of-the-art method [20]. We train our θ -RBM using the Gaussian-Bernoulli formulation in [2], which can be straightforwardly done by adapting Equations (1) and (3,4) such that they can deal with continuous variables. To demonstrate that θ -RBM does not need the input data to span the space of transformations in another experiment we train the proposed model with *mnist* dataset and test it with *mnist-rot*. Since our method relies on an estimation of the dominant orientation of an input image, we perturbed the estimate of orientation and evaluate classification accuracy to test robustness.

Data normalisation: Before training the networks, we normalised the data such that they have zero mean and unit variance for each data component (i.e., the same input feature –pixel location). In case of θ -RBM, normalisation is performed within data having the same orientation.

Discriminative power of learned features: By these experiment we want to demonstrate that our method is able to learn rotation-invariant features, which have a strong discriminative power. We trained the networks using the following parameters: 500 of hidden units, learning rate

Table 1: Classification of *mnist-rot* digits, using RBF-SVM classifier. We report the testing error, obtained for each method, using different sparsity values.

	Sparsity 0.3	Sparsity 0.2	Sparsity 0.1
RBM	15.41%	16.04%	16.39%
O-RBM	15.61%	15.87%	16.61%
TI-RBM	14.07%	13.35%	14.08%
θ-RBM	10.31%	10.08%	10.85%

Table 2: The γ -score, as described in Section 4, on features learned by different approaches. All the trained models used Gaussian-Bernoulli formulation to treat continuous-value data.

	Sparsity 0.3		Sparsity 0.2		Sparsity 0.1	
	Train	Test	Train	Test	Train	Test
RBM	0.1746	0.1775	0.1724	0.1751	0.1721	0.1746
O-RBM	0.1754	0.1782	0.1737	0.1764	0.1710	0.1735
TI-RBM [20]	0.8434	0.8295	0.8106	0.8118	0.7878	0.7917
θ-RBM	0.9000	0.9062	0.9087	0.9098	0.9103	0.9093

$\eta = 0.01$, constant momentum $\alpha = 0.9$, and $S = 9$ rotations in the support set $\mathcal{S} = \{0^\circ, 40^\circ, 80^\circ, 120^\circ, 160^\circ, 200^\circ, 240^\circ, 280^\circ, 320^\circ\}$ (where applicable). We trained an RBF-SVM classifier [4], where parameters were found via grid search. Best results were obtained with loss parameter $C = 10$ and parameter of the Gaussian kernel $\sigma = 0.02$. In Table 1 we show the testing error on the *mnist-rot* dataset. Observe that our method has the lowest error. In Figure 2 we show a subset of filters that θ -RBM learned on the *mnist-rot* dataset. Observe that filters are rotated versions of each other, giving experimental evidence on the validity of 1.

Demonstrating invariance: By these experiments, we want to assess the degree of invariance that θ -RBM can achieve. We use *mnist-rot* for the empirical demonstration of 1, using the same experimental setup as above. In order to compute the γ -score (c.f. Equation (9)), we need a set of transformations to compute the μ_j (c.f. Equation (8)). Instead of using the same support set used at training time, we instead used $\mathcal{S}' = \{\varphi_s + \Delta : \varphi_s \in \mathcal{S}\}$ to compute μ_j , that is a set of rotations that are derived from \mathcal{S} , but are shifted by Δ . We do that because it would result on extremely high invariance score, since the set of rotations is the one that the algorithms were trained with. In our experiments, we set $\Delta = 20^\circ$. Table 2 we show the γ -score for our method, baseline RBM and O-RBM, and TI-RBM, applying different sparsity target. Overall our method achieves the highest γ -score, compared with the other methods. Different sparsity values appears not to influence our results, which is in stark contrast to what happens for TI-RBM, which has -5% loss when the sparsity target decreases. With the achieve γ -score of ~ 0.9 by our method even empirically we can see the validity of 1.

Training on *mnist*: We trained our network using the *mnist* dataset, which does not contain rotated digit images. Since the *mnist-rot* training set has 10,000 images, to achieve equal datasets and meaningful comparisons, we randomly sample the *mnist* to have 10,000 training data as well. Keeping all the parameters the same and for sparsity 0.3, experimental results showed that on the training set our θ -RBM had a lower γ -score (~ 0.77) compared to TI-RBM (~ 0.82). However, TI-RBM had a much lower γ -score on the testing set ~ 0.25 , compared with our method (~ 0.77), highlighting the fact that TI-RBM tends to overfit across the support set of transformations.

Effect of orientation estimation: With this experiment, we want to demonstrate that even when the estimation of the dominant orientation is imperfect our system is still able to learn discriminative rotation-invariant features. To simulate this we randomly perturb the estimated orientation by adding an $\pm\epsilon$, which is an error term drawn from a certain probability distribution. Namely, since orientation is enumerated in \mathcal{S} with an index s , that is if an image x has orientation of $\varphi_s \in \mathcal{S}$, hence it has the s -th orientation, the ϵ affects this assignment, such that $s \leftarrow s \pm \epsilon$. In our experiments we wanted to have an ϵ of up to $n \in \mathbb{Z}$ with a certain probability p . Therefore we used a Bernoulli probability, such that $\epsilon \sim B(n, p)$, which indicates that ϵ can be up to n with a probability of success p .³ Experimental

³The sign of the error term can be drawn from an uniform distribution.

Table 3: Testing error on *mnist-rot* dataset. The training was done on biased estimation of the image orientation with an error drawn from a Bernoulli distribution $B(n, p)$. We reported the testing error, using different values for n and p , namely introducing more error on the estimation of the orientation. For explanation on the * see text.

	$p = 0.1$	$p = 0.2$	$p = 0.3$	$p = 0.4$
$n = 1$	8.9%	9.1%	9.0%	9.6%
$n = 2$	8.9%	9.4%	9.3%	11.0%*
$n = 3$	9.0%	9.9%	10.0%	13.4%*
$n = 4$	9.0%	10.9%*	11.3%*	16.1%*

settings remained the same as previously with sparsity 0.3. In Table 3 we show the experimental results, in terms of testing error, using a RBF-SVM classifier. The SVM has trained using the same setup as reported above. Comparing these results with Table 1, we see that even with large probability of error and even up to 4 errors (in index location) the testing error is not affected (in Table 3 the * denotes the values that are larger than the unperturbed value of 10.31 in Table 1).

Implementation details: We used our own implementation of Restricted Boltzmann machine, which we then used to build upon our θ -RBM. Our code was written in MATLAB and runs on a machine with CUDA capabilities (NVIDIA Titan X). We download TI-RBM from the following URL: https://github.com/kihyuks/icml2012_tirbm, in order to evaluate their method with ours. Our θ -RBM can be downloaded by running the following command: `git clone https://bitbucket.org/stsaft/rbm.git`.

6 Conclusions

In this paper we proposed an unfactored gated restricted Boltzmann machine to learn rotation-invariant features, which were demonstrated to also have strong discriminative power. Differently than the state of the art, our θ -RBM does not need any dataset augmentation. Moreover, it does not even require transformed versions of input images, as it is done e.g. in [12]. The network takes as input an image and its dominant orientation, which is computed by means of histogram of gradients. The third-order tensor, which connects the three layers in the network, is updated using an extended version of the Contrastive Divergence (CD) algorithm [7]. Since each image is associated to a particular slice in the tensor, determined by its dominant orientation, the gradient of the slice is computed using the original CD. Then, the remaining slices are generated by rotated versions of the computed gradient.

Using an SVM classifier, our method had the lowest classification error, compared with the current state of the art, as shown in Table 1. Moreover, we mathematically proved that our method learns rotation invariant features, which was also demonstrated experimentally in Table 2, using the invariance measure proposed in [16], over the *mnist-rot* dataset [11]. This measure ranges in $[0, 1]$, where ‘1’ means that features are fully invariant to transformations (in our case, rotations). In fact, our proof shows that our learning algorithm can reach the highest value theoretically. Furthermore, experimentally our method had an invariance score of $\sim 90\%$, which is the highest amongst baseline approaches and current state of the art approaches in transformation invariant RBM [20].

Acknowledgements

We thank NVIDIA corporation for providing us a GPU Titan X video card.

References

- [1] D. Cheng, T. Sun, X. Jiang, and S. Wang. Unsupervised feature learning using Markov deep belief network. In *2013 IEEE ICIP*, pages 260–264. IEEE, Sep 2013. ISBN 978-1-4799-2341-0.
- [2] K. Cho, A. Ilin, and T. Raiko. Improved learning of Gaussian-Bernoulli restricted Boltzmann machines. *LNCS*, 6791:10–17, 2011. ISSN 03029743.
- [3] A. Coates, A. Arbor, and A. Y. Ng. An Analysis of Single-Layer Networks in Unsupervised Feature Learning. *AISTATS*, pages 215–223, 2011.

- [4] C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995.
- [5] A. Dosovitskiy, J. T. Springenberg, M. Riedmiller, and T. Brox. Discriminative Unsupervised Feature Learning with Convolutional Neural Networks. *NIPS* 27, pages 1–13, 2014. ISSN 0162-8828.
- [6] I. J. Goodfellow, Q. V. Le, A. M. Saxe, H. Lee, and A. Y. Ng. Measuring Invariances in Deep Networks. *NIPS* 22, 22:646–654, 2009.
- [7] G. E. Hinton. Training products of experts by minimizing contrastive divergence. *Neural computation*, 14(8):1771–1800, 2002. ISSN 0899-7667.
- [8] G. E. Hinton, S. Osindero, and Y.-W. Teh. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–54, 2006. ISSN 0899-7667.
- [9] M. Jaderberg, K. Simonyan, A. Zisserman, and k. Kavukcuoglu. Spatial transformer networks. In *NIPS* 28, pages 2017–2025. Curran Associates, Inc., 2015.
- [10] N. Karianakis, Y. Wang, and S. Soatto. Learning to discriminate in the wild: Representation-learning network for nuisance-invariant image comparison. Technical report, UCLA Vision Lab, 2013.
- [11] H. Larochelle, D. Erhan, A. Courville, J. Bergstra, and Y. Bengio. An empirical evaluation of deep architectures on problems with many factors of variation. In *ICML*, pages 473–480. ACM, 2007. ISBN 978-1-59593-793-3.
- [12] R. Memisevic. Learning to relate images. *PAMI*, 35(8):1829–1846, 2013. ISSN 01628828.
- [13] R. Memisevic and G. Hinton. Unsupervised learning of image transformations. In *CVPR*, pages 1–8, June 2007.
- [14] R. Memisevic and G. Hinton. Learning to Represent Spatial Transformations with Factored Higher-Order Boltzmann Machines. 1492:1473–1492, 2009.
- [15] D. C. Mocanu, H. Bou Ammar, D. Lowet, K. Driessens, A. Liotta, G. Weiss, and K. Tuyls. Factored four way conditional restricted Boltzmann machines for activity recognition. *Pattern Recognition Letters*, 66:100–108, 2015. ISSN 01678655.
- [16] A. Rasmus, T. Raiko, and H. Valpola. Denoising autoencoder with modulated lateral connections learns invariant representations of natural images. (2014):1–15, 2015. URL <http://arxiv.org/abs/1412.7210>.
- [17] U. Schmidt and S. Roth. Learning rotation-aware features: From invariant priors to equivariant descriptors. In *CVPR*, pages 2050–2057, 2012. ISBN 9781467312264.
- [18] Z. Shou, Y. Zhang, and H. J. Cai. A study of transformation-invariances of deep belief networks. In *IJCNN*, pages 1–8. IEEE, aug 2013. ISBN 978-1-4673-6129-3.
- [19] P. Smolensky. Information processing in dynamical systems: Foundations of harmony theory. *Parallel Distributed Processing Explorations in the Microstructure of Cognition*, 1(1):194–281, 1986.
- [20] K. Sohn and H. Lee. Learning Invariant Representations with Local Transformations. *ICML*, pages 1311–1318, 2012.
- [21] K. Sohn, G. Zhou, C. Lee, and H. Lee. Learning and Selecting Features Jointly with Point-wise Gated Boltzmann Machines. *ICML*, 28, 2013.
- [22] J. Susskind, G. Hinton, R. Memisevic, and M. Pollefeys. Modeling the joint density of two images under a variety of transformations. *CVPR*, pages 2793–2800, 2011. ISSN 10636919.
- [23] W. Zou, S. Zhu, K. Yu, and A. Y. Ng. Deep learning of invariant features via simulated fixations in video. In *NIPS* 25, pages 3203–3211. Curran Associates, Inc., 2012.